

OCT 05 2006

AMENDMENTS TO THE CLAIMS:

1. (Original) A method of reading data comprising the steps of:
receiving a request for a stripe of erasure coded data stored across a plurality of storage devices, the stripe comprising stripe blocks;
sending read messages to at least a quorum of the storage devices;
receiving at least the quorum of reply messages from the storage devices, the quorum of the reply messages including at least a minimum number of the stripe blocks needed to decode the stripe of erasure coded data, the quorum meeting a quorum condition of a number such that any two selections of the number of the stripe blocks intersect in the minimum number of the stripe blocks; and
decoding the stripe of erasure coded data from at least the minimum number of the stripe blocks, thereby forming the data.
2. (Original) The method of claim 1 wherein each of the reply messages within the quorum indicates that there is no pending write for the stripe block stored on the storage device associated with the reply message.
3. (Original) The method of claim 1 wherein each of the reply messages within the quorum indicates that the stripe block associated with the reply message has a timestamp that matches other timestamps associated with other reply messages within the quorum.

4. (Original) The method of claim 1 wherein a coordinator performs the steps.

5. (Original) The method of claim 4 wherein the coordinator comprises one of the storage devices.

6. (Original) The method of claim 5 wherein the coordinator effectively sends one of the read messages to itself.

7. (Original) The method of claim 6 wherein the coordinator effectively receives one of the reply messages from itself.

8. (Original) A method of reading data comprising the steps of:

receiving a request for a stripe of erasure coded data stored across a plurality of storage devices, the stripe comprising stripe blocks which comprise a first number of data blocks and a second number of parity blocks;

sending read messages to the storage devices;

receiving at least a quorum of reply messages from the storage devices

which indicate that there is no pending write for the stripe block stored on the storage device, the quorum of the reply messages including at least the first number of the stripe blocks, the quorum comprising the first number plus a half of the second number; and

decoding the stripe of erasure coded data from the first number of the

stripe blocks, thereby forming the first number of the data blocks.

9. (Original) The method of claim 8 wherein the quorum of the reply messages includes validation timestamps which match.

10. (Original) The method of claim 8 wherein a coordinator performs the steps.

11. (Original) The method of claim 10 wherein the coordinator comprises one of the storage devices.

12. (Original) The method of claim 11 wherein the coordinator effectively sends one of the read messages to itself.

13. (Original) The method of claim 12 wherein the coordinator effectively receives one of the reply messages from itself.

14. (Original) The method of claim 8 wherein the coordinator is not one of the storage devices upon which the stripe of erasure coded data is stored.

15. (Original) The method of claim 8 further comprising the step of identifying a group of the storage devices as targets.

16. (Original) The method of claim 15 wherein the step of identifying the targets randomly picks the targets.

17. (Original) The method of claim 15 wherein each of the query messages sent to the targets identifies the storage device as one of the targets.

18. (Original) The method of claim 17 wherein the reply messages from the targets include the stripe blocks.

19. (Original) The method of claim 8 wherein the storage devices comprise a distributed storage system.

20. (Original) The method of claim 19 wherein the distributed storage system comprises a quantity of the storage devices.

21. (Original) The method of claim 20 wherein the quantity of the storage devices corresponds to the first number of the data blocks plus the second number of the parity blocks.

22. (Original) The method of claim 20 wherein the quantity of the storage devices exceeds the first number of the data blocks plus the second number of the parity blocks.

23. (Original) The method of claim 22 further comprising the step of identifying the storage devices upon which the stripe of erasure coded data is stored.

24. (Original) The method of claim 8 wherein each of the storage devices comprises a log, wherein the log comprises log entries of each successful write of data, the log entries comprising a stripe indicator, a write timestamp, and a physical location of the stripe block on the storage device.

25. (Original) The method of claim 8 wherein the read messages include a stripe indicator.

26. (Original) The method of claim 8 wherein the reply messages include a stripe indicator.

Cancel Claims 27-28.

29. (Original) A computer readable memory comprising computer code for implementing a method of reading a stripe of erasure coded data, the method of reading the stripe of erasure coded data comprising the steps of:
receiving a request for a stripe of erasure coded data stored across a plurality of storage devices, the stripe comprising stripe blocks;
sending read messages to at least a quorum of the storage devices;
receiving at least the quorum of reply messages from the storage

devices, the quorum of the reply messages including at least a minimum number of the stripe blocks needed to decode the stripe of erasure coded data, the quorum meeting a quorum condition of a number such that any two selections of the number of the stripe blocks intersect in the minimum number of the stripe blocks; and decoding the stripe of erasure coded data from at least the minimum number of the stripe blocks, thereby forming the data.

30. (Original) The computer readable memory of claim 29 wherein each of the reply messages within the quorum indicate that there is no pending write for the stripe block stored on the storage device associated with the reply message.

31. (Original) The computer readable memory of claim 29 wherein each of the reply messages within the quorum indicate that the stripe block associated with the reply message has a timestamp that matches other timestamps associated with other reply messages within the quorum.

Cancel Claim 32.